Declining Volatility in the U.S. Automobile Industry

By

Valerie A. Ramey

University of California, San Diego and National Bureau of Economic Research

and

Daniel J. Vine

Board of Governors of the Federal Reserve System

August 2005

We are indebted to Edward Cho and Jillian Medeiros for outstanding research assistance and to George Hall for his Chrysler data. We have benefited from helpful comments from Garey Ramey. Valerie Ramey gratefully acknowledges support from National Science Foundation grant #0213089.

1 Introduction

The U.S. automobile industry is an important source of volatility in the U.S. economy. In the last 40 years, motor vehicle production has accounted for almost 25 percent of the variance in aggregate GDP growth, even though gross motor vehicle output represents less than five percent of the *level* of aggregate GDP.[1] In the mid-1980s, however, the variance of automobile production declined drastically, falling by more than 70 percent relative to its past level. While the variance of auto sales also receded, the decline was smaller than the decline in output volatility. Moreover, the covariance of inventory investment with sales became negative, suggesting that automakers began to more actively use inventories to insulate production from sales shocks in the 1980s. At the same time, fluctuations in production at U.S. assembly plants, which in the 1970s and early 1980s largely reflected changes in the number of workers attached to each plant, began instead to reflect variations in average hours per worker in the 1990s. Interestingly, a number of these changes were observed at the aggregate level as well.[2]

This paper documents these developments in the U.S. auto industry, and, building on the work of Blanchard (1983), shows how the changes observed in sales, inventories and production behavior in the 1980s could have stemmed from one underlying factor—a decline in the persistence of motor vehicle sales. We use both industry-level data as well as micro data on

---

[1] The data appendix gives details of all calculations and estimates.

[2] Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) document a 50 percent decline in the variance of GDP growth beginning in 1984. Golob (2000) shows that the covariance of inventory investment and sales switched from being positive in most industries before 1984 to being negative after 1984. Kahn, McConnell and Perez-Quiros (2002) showed that within durable goods manufacturing the variance of production fell much more than the variance of sales.

production schedules from 103 assembly plants in the United States and Canada to document the developments in the early 1980s. Analyzing the original Holt, Modigliani, Muth and Simon (1960) version of the linear-quadratic inventory model, we show that a decline in the persistence of sales leads to all of the changes noted above, even in the absence of technological change.

2 Structural Change in the U.S. Automobile Industry: The Facts

2.1 Production, Sales, and Inventory Variances

The history of seasonally-adjusted car and truck production (in physical units) is shown in Figure 1 for the months of January 1967 through December 2004. Also shown are seasonally-adjusted sales for domestically-produced vehicles in each of these market segments, which include vehicles assembled in the U.S., Canada and Mexico. The truck market segment includes vans and SUVs.

The graphs in Figure 1 clearly show that the secular trend in the car segment, where sales and production have been relatively flat over history and have declined in recent years, differs notably from the trend in the truck segment, where production and sales have increased steadily since 1980. Since there are obvious differences in the conditional means of these two market segments, we treat them separately in most of the analysis below.

The variances of both car and truck output (relative to their respective trends) dropped sharply in the mid-1980s. According to structural break tests on the variance of detrended seasonally-adjusted car production, there was a statistically significant break between February and March in 1984. For truck production, the break lies between January and February in 1983.[3] Thus, the structural break in the variance of automobile production occurs at essentially the same time as the structural break in GDP growth volatility, which many studies place in 1984.

In order to quantify the change in industry volatility, we measure the variances of key variables in the auto industry in the two sample periods divided by 1984. The variables of interest are derived from the standard inventory identity: $Y_t = S_t + \Delta I_t$, where Y is production, S is sales, and ΔI is the change in inventories. For stationary variables, the following relationship exists between the variance of production and the variance of sales:

$$Var(Y_t) = Var(S_t) + Var(\Delta I_t) + 2Cov(S_t, \Delta I_t). \tag{1.1}$$

Table 1 reports this variance decomposition for cars and trucks. Because production and sales have trends within each sample, production, sales and inventory investment are first normalized by their respective trends. The variance of production for cars fell by 70 percent after 1984; for trucks, it fell by 87 percent. Moreover, for both cars and trucks the variance of production fell by a larger percentage than sales, and the covariance of inventory investment with final sales either switched from being positive to being negative or became more negative.

The variances and covariances do not add up in part because domestic sales and inventories include imports from Canada and Mexico, and a small portion of U.S. production is

[3] One important difference between the time-series properties of physical unit data and NIPA data used in other studies is that stationarity tests on the logarithm of physical unit variables reject a unit root in favor of a deterministic trend. To search for the structural break in the variance, we used data detrended with an HP filter rather than trend breaks so as not to bias the results for a particular period. In particular, we used seasonally adjusted output divided by the exponential of the HP filter trend applied to the log of output. The p-values were essentially zero.

exported.[4] Changes in North American motor vehicle trade behavior, however, do not appear to explain the structural change in industry volatility. The addendum to Table 1 shows the trade-augmented variance decomposition for cars, where U.S. production is augmented with North American imports and U.S. sales are augmented with exports. The changes in the variances after 1984 are very similar to those shown in Table 1. (Trade data for trucks are not readily available.)

Golob (2000) and Kahn, McConnell, Perez-Quiros (2002) uncovered a similar decline in the volatility of aggregate chain-weighted durable goods output. Some researchers have linked this decrease in volatility to a decline in the inventory-sales ratio. However, as shown in Figure 2, the inventory-sales ratio ("days-supply") for light vehicles shows no evidence of change after 1984.[5] Thus, the changes we have documented in the auto industry occurred for some other reason.[6]

2.2 Intensive and Extensive Labor Margins

Short-run fluctuations in the number of assemblies traditionally come from automakers changing either their work schedule or the rate of production in their assembly plants.[7] Changes

---

[4] In addition to trade, the variance identity does not hold exactly because each variable is seasonally adjusted and detrended separately.

[5] Because data for inventories of light trucks are not available before 1972, the pre-1972 numbers are for cars only. During the early 1970s, the inventory-sales ratio for cars and trucks were very similar.

[6] This does not imply that improvements in inventory management had no impact on the auto industry. In fact, the auto industry pioneered *just-in-time* inventories in the 1980s, and the ratio of materials and work-in-progress inventories to shipments for the entire automobile industry (SIC code 371) did fall after 1984. That decline cannot, however, explain the changes we document for the relationship between production, sales, and inventories of finished vehicles.

[7] Plant entry and exit historically are responsible for the secular trend in output. (See Bresnahan and Ramey (1984).)

to the plant schedule may be temporary, such as adding overtime hours or closing a plant for a week (called an "inventory adjustment"), or persistent, such as adding a second shift. Automakers change the rate of production at a plant by raising or lowering the line speed, which has historically required a change in the number of workers on each shift.

While fluctuations in total worker hours depend mostly on the number of vehicles assembled, fluctuations in employment and average hours depend on the type of output adjustments that are chosen. Overtime hours and inventory shutdowns are intensive adjustments to hours worked, while changes to the number of shifts and changes to the line speed are largely extensive in nature. Although inventory adjustments involve temporary layoffs, they typically last only a week and involve negligible adjustment costs. They essentially serve as a means to change the average hours per worker measured on a monthly basis. The costs associated with using these margins have often been named as the major source of volatility in the auto industry (see Aizcorbe (1992), Bresnahan and Ramey (1994) and Hall (2000)). To summarize several key results, adjustments to production on the extensive margins entail large adjustments costs and are therefore used when a demand change is perceived to be persistent in nature. Transitory shocks to sales, on the other hand, are best handled with adjustments to the intensive margins. While overtime hours and inventory adjustments do affect marginal costs, their use incurs no adjustment costs.

Table 2 measures the contribution of intensive and extensive adjustments in the overall variance of monthly output using a dataset that tracks production schedules at U.S. and Canadian assembly plants operated by the domestic automakers. Separate measurements are made for the pre-1984 and post-1984 periods.[8] Contributions to variance are measured as in Bresnahan and Ramey (1994), where the variance of actual output is compared to an artificial output measure that holds each margin, in turn, constant at each plant. The difference in the variances of actual and constructed output after holidays, supply disruptions, and annual summer shutdowns are removed determines the impact of each margin on the variance of plant-level output. The numbers in Table 2 are weighted averages across all plants and do not sum to 100 because of nonlinearities and covariance terms.

The contribution of adjustments to extensive margins in the variance of output declined from 65 percent in the first period to 37 percent in the second period. A change in the way shifts were added and pared between the two periods accounted for most of this change. The contribution to variance of adding and cutting shifts fell from 33 percent of the monthly production variance in the 1970s and 1980s to only 5 percent during the 1990s.

Adjustments to intensive margins, on the other hand, became more important in the latter period. The contribution of overtime hours and inventory adjustments rose from about 40 percent of plant-level variance in the early period to 51 percent in the latter period. In particular, the contribution of overtime hours to total output variance stepped up significantly from a contribution of less than 10 percent in the early period to more than 24 percent in the second period. Plant closures for inventory adjustment were unchanged.

The increase in the use of overtime hours during the 1990s expansion has been noted in other industries as well (e.g. Hetrik (2000)). One possible explanation for the increase in overtime hours is the rise in the cost of health insurance benefits, which boosts the fixed cost of

---

[8] Our post-1984 data starts in 1990 because we did not have access to *Automotive News* for 1984-1989.

employing a worker. While this change may result in higher average hours per worker, the change in the relative contributions of the intensive versus extensive margins to the variance of total hours is not clearly linked to health insurance costs.

2.3 The Persistence of Sales

The decline in the variance of sales shown in Table 1, though smaller than the decline in the variance of output, could have stemmed from a reduction in the size of sales shocks or from a reduction in their persistence. If sales are represented by a simple AR(1) with a first-order autocorrelation of ρ and an error-term variance of σ^2, for example, their variance is given by $\sigma^2 / (1-\rho^2)$. Most of the reduction in variance of sales during the 1980s, it turns out, came from a weakening in the persistence of sales. This is important because simple inter-temporal models of production scheduling yield solution paths in which the variance of output relative to sales depends crucially on the persistence of sales.

Consider the following univariate model for domestic auto sales between January 1967 and December 2004, where $\tilde{S}$ is the logarithm of seasonally-adjusted sales:

(1.2) $$\tilde{S}_t = \alpha_0 + \alpha_1 \cdot t + \alpha_2 \cdot \tilde{S}_{t-1} + D_t \cdot \left[\beta_0 + \beta_1 \cdot t + \beta_2 \cdot \tilde{S}_{t-1} \right] + \varepsilon_t$$

where

$$\varepsilon_t \sim N\left(0, \sigma^2 + \beta_3 \cdot D_t\right)$$

and

$$D_t = 0 \ \text{ for } t < 1984:1$$
$$D_t = 1 \ \text{ for } t \geq 1984:1.$$

This model allows all parameters of the sales process to change in 1984, including the coefficient on lagged sales, the constant, the slope of the trend, and the variance of the residual. We estimate this model via maximum likelihood for cars alone, light trucks alone, and for the combination of cars and light trucks, called "light vehicles."

The coefficient estimates shown in Table 3 indicate significant changes have occurred in the process governing sales. The constant and the trend differ across the two periods for all three aggregates, which reflects the shift in sales from cars to light trucks shown in Figure 1. The parameter β_3, which measures the change in the variance of the sales shocks, shows a significant decline in 1984 for light trucks but was unchanged for cars and the light-vehicle aggregate. The first-order autocorrelation of sales, on the other hand, fell between the early and the late periods for all three aggregates, as evidenced by the negative and significant point estimates of β_2. For cars the first-order autocorrelation fell from 0.85 to 0.56, and for trucks it fell from 0.93 to 0.69. When all light vehicles are grouped together, this estimate declined from almost 0.9 to 0.6.

The sales process in the post-1984 period returns to its mean much more quickly following a surprise than was the case in earlier decades. It is also clear that most of the change in the unconditional variance of sales described in Table 1 came from a change in the propagation of sales rather than from a change in the variance of sales shocks.

The persistence of sales could have changed for a number of reasons. One possibility is that the automakers began responding to shocks more aggressively with their pricing policies. Another possibility is that the types of shocks hitting the industry, such as oil price shocks, became less persistent. The production-scheduling model that we present in the next section does not depend on the source of the change in persistence, but only on the fact that it occurred.

We take this change in persistence as a given and examine the implications for the behavior of production, inventories, average hours and employment.

3 Insights from the Holt, Modigliani, Muth and Simon Model of Production Scheduling

The changes in the sales process described in the preceding section have large effects on the relationship between production, inventories and sales. The importance of the persistence of sales in production volatility in the auto industry was, in fact, acknowledged in an inventory study by Blanchard (1983) before any changes in auto sales and production had occurred. Blanchard concluded that costs in the auto industry were such that inventories could either stabilize or destabilize production depending on the persistence of sales. Since fluctuations in monthly automobile sales were very persistent during Blanchard's sample period, 1966 – 1979, inventories destabilized production.

In Ramey and Vine (2004), we analyzed a model of industry costs that accounted for the nonconvexities and lumpiness in production margins that are a key feature of automobile assembly. In this paper, we analyze the Holt, Modigliani, Muth and Simon (1960) model of industry costs. While this convex cost model is a less accurate description of the automobile industry, it does capture the key distinctions of the intensive and extensive margins and has the advantage of conveying the intuition more clearly, of being solved analytically, and of having broader applicability to other industries. Moreover, the basic results concerning the effects of sales persistence changes are the same in both types of models.

In this section we show how the persistence of sales changes the decision rules for inventories and workforce in a production-smoothing model. In the next section, we show how the persistence of sales affects the variances of these variables.

3.1 Production-Scheduling Model with Inventories and Workforce

The original production smoothing model of inventories was introduced by Holt, Modigliani, Muth and Simon in 1960. The model uses quadratic approximations to the various costs faced by the factory manager.

Consider the problem of a firm that faces a stochastic sales process and must choose the size of its workforce, N_t, in addition to the level of output, Y_t, each week to minimize the discounted present value of production, workforce-adjustment and inventory-holding costs. The choice of workforce determines the firm's minimum efficient scale of production in each period— θN. The cost-minimization problem takes the form of Equation (2.1), where I_t is the stock of inventories at the end of period t, S_t is sales in period t, and ε_t is an i.i.d. shock to sales:

$$(2.1)\quad \underset{\{N_{t+j},Y_{t+j}\}}{\mathrm{MIN}}\, C_t = E_t\left\{\lim_{J\to\infty}\sum_{j=0}^{J}\beta^{t+j}\cdot\frac{1}{2}\left[\gamma_1 Y_{t+j}+\gamma_2(Y_{t+j}-\theta N_{t+j})^2+\gamma_3 N_{t+j}+\gamma_4(N_{t+j}-N_{t+j-1})^2+\alpha_1(I_{t+j-1}-\alpha_2 S_{t+j})^2\right]\right\};$$

$$0<\beta<1,\ \gamma_1\geq 0,\ \gamma_2\geq 0,\ \gamma_3\geq 0,\ \gamma_4\geq 0,\ \alpha_1>0,\ \text{and}\ \alpha_2\geq 0.$$

The minimization is subject to the inventory identity,

$$(2.2)\qquad Y_{t+j}=I_{t+j}-I_{t+j-1}+S_{t+j},$$

and the process governing sales,

$$(2.3)\qquad S_{t+j}=c+\rho S_{t+j-1}+\varepsilon_{t+j},$$

for $j\geq 0$ and with $E(\varepsilon_{t+j}) = 0$ and $Var(\varepsilon_{t+j}) = \sigma^2_{\varepsilon_S}$. The firm observes S_t before it chooses employment and production in period t.

The second term in Equation (2.1) captures the cost of scheduling a workweek that either exceeds or falls below what is considered "full-time," where θ is the product of a normal workweek (such as 40 hours). To see this, let $Y_t = h_t \cdot N_t$, where h_t represents average hours per worker, and rewrite this term as $\gamma_2((h_t - 40)N_t)^2$. This term captures the cost per worker of scheduling overtime or short weeks, which is an intensive adjustment.

The fourth term in Equation (2.1) is the cost of adjusting the number of workers attached to the plant—the extensive margin. Adjustments to the workforce size essentially shift the static marginal cost curve horizontally and re-define the minimum efficient scale of production. In contrast to varying the workweek, increasing the number of workers does not lead to increasing static marginal costs but this move does incur dynamic adjustment costs.

The last term captures the trade-off between inventory-holding costs and stock-out costs, which depends on the level of sales. For industries that produce to stock, such as motor vehicles, this is a standard way to obtain an industry equilibrium in which inventories are non-zero.

The modern production-smoothing model of inventory behavior is a simplified version of this original Holt et al model. The models used by Blanchard (1983) and those surveyed in Ramey and West (1999), for example, do not distinguish between the intensive and extensive margins of labor input, and thus are special cases of the Holt et al model in which θ = 0 and output equals (or is in fixed proportion to) workforce.[9] All increases in production imply a rising marginal cost in these models, and there is no distinction between boosting the workweek and hiring more workers. This distinction, however, is very important in the auto industry, as

[9] One reason that the special case of the original inventory model became dominate is that it has only one endogenous state variable—the level of lagged inventories. In contrast, the Holt et al model has an additional endogenous state variable—the level of lagged workforce. Going from one state variable to two state variables makes the system significantly more difficult to solve analytically.

documented by numerous authors (e.g. Aizcorbe (1992), Bresnahan and Ramey (1994), and Hall (2000).

3.2 Solution: Optimal Production Scheduling

For simplicity we set γ_1 , γ_3 and c equal to zero, since these linear terms affect only the means of the variables in the solution and not the dynamics. The first-order conditions with respect to workforce and inventories in the current period are written respectively as:

$$\gamma_4\left(N_t - N_{t-1}\right) - \theta\gamma_2\left(Y_t - \theta N_t\right) = E_t\left\{\beta\gamma_4\left(N_{t+1} - N_t\right)\right\} \tag{2.4}$$

and

$$E_t\left\{\gamma_2\left(Y_t - \theta N_t\right) + \beta\alpha_1\left(I_t - \alpha_2 S_{t+1}\right)\right\} = E_t\left\{\beta\gamma_2\left(Y_{t+1} - \theta N_{t+1}\right)\right\}. \tag{2.5}$$

First-order condition (2.4) states that employment, given some level of output, is optimized when the cost of adding one more worker this period, less the savings in current-period production costs, equals the discounted cost of adjusting workforce by one less worker next period. First-order condition (2.5) is analogous to the first-order condition obtained in the simple model without workforce, and it states that the cost of producing one more unit in the current period and storing it in inventory equals the discounted saving of producing one less unit next period. When workforce and output decisions are optimized simultaneously, the solution path for the system is homogeneous of degree zero in α_1, γ_2 and γ_4.

The decision rules for N_t and I_t, assuming rational expectations, depend on two state variables—N_{t-1}, I_{t-1}—and on S_t, as shown in Equation (2.6).[10]

(2.6)
$$\begin{pmatrix} N_t \\ I_t \end{pmatrix} = C \begin{pmatrix} N_{t-1} \\ I_{t-1} \end{pmatrix} + \vec{d} \cdot S_t$$

The persistence of sales, ρ, affects only the coefficients in $\vec{d}$ and not those in $\boldsymbol{C}$.

4 Variances of Output, Inventories and Workforce

Expressions for the variances of output, inventories and workforce and the covariance between inventory investment and sales are readily derived from the decision rules. Because the decision rules from the general Holt et al model are rather complicated functions of the parameters, however, it is first useful to analyze two special cases in which it is easier to see the intuition of how sales persistence affects these moments. The results are then shown to hold up in the general model as well.

4.1 Sales Persistence and the Variance of Output: The Simplified Modern Production Smoothing Model

Consider a special case of the Holt et al model that sets $\theta = 0$ and $\gamma_4 = 0$. This is a simpler version of the model used by Blanchard (1983), in which $\theta = 0$ and $Y = N$. For a particular parameterization of his model, he showed that the production response to a sales shock was greater if the persistence of sales was greater.

---

[10] See the Mathematical appendix for the full solution.

The optimal decision rule for production in this simplified model, assuming rational expectations, is given by:

$$Y_t = -(1-\lambda) I_{t-1} + \phi S_t \,, \tag{2.7}$$

where

$$\lambda = \frac{1}{2}\left\{ \frac{1}{\beta} + 1 + \frac{\alpha_1}{\gamma_2} - \sqrt{\left[\frac{1}{\beta} + 1 + \frac{\alpha_1}{\gamma_2}\right]^2 - \frac{4}{\beta}} \right\}$$

and

$$\phi = \frac{1 - \lambda + \beta\lambda\rho \dfrac{\alpha_1\alpha_2}{\gamma_2}}{1 - \beta\lambda\rho} \,.$$

Note that even without adjustment cost on production (or employment), the model is nevertheless dynamic because the convex production cost operates like an adjustment cost on inventories, and the change in inventories is a function of production.

As long as the α*'s* and γ*'s* are nonnegative, λ will lie between zero and unity, and ϕ will be positive. While λ depends on neither α_2 (the desired inventory-sales ratio) nor ρ (the persistence parameter for sales shocks), ϕ is increasing in both of these parameters and will play an important roll below.

The relative variances of production and sales and the covariance of sales and inventory investment in the simplified model are given by the following expressions:

$$\frac{\sigma_Y^2}{\sigma_S^2} = 1 + \frac{2(1-\rho)(\phi-1)}{1-\lambda\rho} + \frac{2(\phi-1)^2(1-\rho+\rho\lambda^2)}{(1+\lambda)(1-\lambda\rho)} \tag{2.8}$$

and

(2.9) $$\sigma_{S,\Delta I} = \frac{(\phi - 1)\sigma^2_{\varepsilon_S}}{(1-\lambda\rho)(1+\rho)}.$$

The value of ϕ relative to unity is an important determinant of the relationship between production, sales and inventories. If $\phi > 1$ the covariance is positive and the variance of production is greater than the variance of sales. If $\phi < 1$ the covariance of sales and inventory investment is negative, and the variance of production is potentially, but not necessarily, greater than the variance of sales. The covariance is more likely to be positive when shocks to sales are persistent, however, as ϕ is an increasing function in ρ.

An increase in either α_2 or ρ generally raises the variance of output relative to sales. If a firm maintains a higher ratio of inventories to sales (higher α_2) then a given increase in sales will lead the firm to produce more. The size of the output response depends on the cost of deviating from the desired ratio relative to the marginal cost of production.

If ρ is close to one, the firm anticipates that sales will remain elevated for a long time following a positive shock, so it raises production in order to prevent its inventory-sales ratio from dipping too low for an extended period. If ρ is close to zero, on the other hand, the sales shock is temporary and it is cheaper to let inventories stray from the desired inventory-sales ratio for a few periods rather than boost output by the full amount necessary to accommodate sales and maintain the desired inventory-sales ratio concurrently. In the numerator of ϕ, ρ and α_2 multiply each other, so there is also an interaction effect.

Thus, a change in ρ alone can lead to a change in signs of the covariance of sales and inventory investment from positive to negative. Neither of these expressions is monotonically increasing in ρ for all possible parameter values, however. In order to determined sufficient conditions under which these measures are strictly increasing in ρ, we numerically investigated the parameter space for these functions, with β pre-set to 0.997 (annual discount rate of four percent). We searched over values of α_2 from 0 to 5 months, values of α_1/γ_2 from 0.001 to 5, and values of ρ from 0.01 to 0.99.

For virtually all parameter values, the covariance of sales and inventory investment is monotonically increasing in ρ. The derivative of the covariance with respect to ρ is positive for any value of α_2 as long as $\alpha_1/\gamma_2 > 0.05$, so a rise in ρ always leads to an increase in the covariance of sales and inventory investment as long as the penalty for deviating from desired inventories is not too small relative to the marginal cost of production. For $\alpha_2 = 2.5$, which coincides with the average inventory-sales ratio in the automobile industry (stated in months), $\alpha_1/\gamma_2 > 0.027$ guarantees that the covariance is increasing in ρ.

Most parameter values also imply that an increase in ρ leads to an increase in the variance of production relative to the variance of sales. For $\rho < 0.6$, σ_Y^2/σ_S^2 is monotonically increasing in ρ for all values of α_2 and α_1/γ_2 within the ranges explored. Figure 3 shows the regions of the parameter space where the derivative is positive for various values of ρ. As ρ rises, the part of the parameter space for which the derivative is positive shrinks.[11]

---

[11] The reason for this is most easily seen by decomposing the variance of output as in Equation (1.1) and dividing by the variance of sales: $\sigma_Y^2/\sigma_S^2 = 1 + \sigma_{\Delta I}^2/\sigma_S^2 + 2\cdot\sigma_{S,\Delta I}/\sigma_S^2$. The variances σ_Y^2 and σ_S^2 increase to infinity as $\rho \to 1$ while $\sigma_{\Delta I}^2$ and $\sigma_{S,\Delta I}$ remain finite. The second and third terms thus shrink to zero and the expression returns to 1 in the limit.

4.2 Sales Persistence, Hours, and Employment: The Case with No Inventories

Consider next the special case in which firms cannot hold inventories but are able to choose workforce (ie. $Y_t = S_t$ and $\alpha_1 = \alpha_2 = 0$). The stochastic Euler equation for this problem is given by:

$$E_t\left\{\beta\gamma_4 N_{t+1} - \left[(1+\beta)\gamma_4 + \gamma_2\theta^2\right]N_t + \gamma_4 N_{t-1}\right\} = -\gamma_2\theta S_t \tag{3.1}$$

The rational expectations solution to equation (3.1) that satisfies the transversality condition is:

$$N_t = \lambda N_{t-1} + \frac{\gamma_2\theta}{\gamma_4}\frac{\lambda}{1-\beta\lambda\rho}S_t \tag{3.2}$$

with

$$\lambda = \frac{1}{2}\left\{\frac{1}{\beta} + 1 + \frac{\gamma_2\theta^2}{\beta\gamma_4} - \sqrt{\left[\frac{1}{\beta} + 1 + \frac{\gamma_2\theta^2}{\beta\gamma_4}\right]^2 - \frac{4}{\beta}}\right\}. \tag{3.3}$$

As long as all parameters are positive, λ will be between zero and unity. The variance of workforce relative to the variance of sales (which in this simple case is equal to the variance of production) depends on the persistence of sales, ρ, as shown in Equation (3.4).

$$\frac{\sigma_N^2}{\sigma_S^2} = \frac{\sigma_N^2}{\sigma_Y^2} = \frac{1}{1-\lambda^2}\left[\frac{\gamma_2\theta}{\gamma_4}\frac{\lambda}{1-\beta\lambda\rho}\right]^2\left[1 + 2\frac{\lambda}{1-\lambda\rho}\right]. \tag{3.4}$$

This expression is always increasing in ρ. The intuition is as follows: Because workforce changes entail an adjustment cost while overtime hours or short weeks do not, workforce accounts for a larger part of output movements only when adjustment costs pay off—when the new workforce size is expected to yield lower static marginal costs well into the future.

Thus, a rise in persistence of sales increases the contribution of workforce to the variance of output.

4.3 The Persistence of Sales and the Variances of Output, Inventories and Workforce in the General Model

In the general model both output and workforce are chosen in each period. While closed-form solutions to the model do exist, the effects of individual parameters on the decision rules are difficult to see in these complicated expressions. Fortunately, the intuition developed in the special cases is largely unchanged when inventories and workforce are optimized jointly.

The variance of output relative to sales, the covariance between sales and inventory investment and the variance of workforce relative to output are plotted against values of ρ in the top row of Figure 4. The solid lines represent these key measures of volatility for a baseline set of parameters chosen so that the volatility of output relative to sales and the volatility of workforce relative to output from the model are consistent with the empirical counterparts from the auto industry prior to 1984.[12] We do not attempt to estimate the parameters from this model because, while having the advantage of being simple and intuitive, the model does not capture the important nonconvexities in the automobile industry.[13] The other (non-solid) lines plot these measures for alternative parameterizations of the model and will be discussed below. The

[12] Specifically, α_2 is set to 2.5, which is the average inventory-sales ratio (in months), θ is normalized to 1, ρ is estimated from the first-order autocorrelation of sales to be 0.85, and β is pre-set to 0.997 (4 percent annual rate). With these parameters in place, $\alpha_1/\gamma_2 = .085$ and $\gamma_4/\gamma_2 = 1.65$ together yield decision rules in which $\sigma_Y^2/\sigma_S^2 = 1.28$ and $\sigma_{\theta N}^2/\sigma_Y^2 = 0.647$; values which match their empirical counterparts for cars in Tables 1 and 2.

[13] Moreover, the worker-output ratio (θ) is not constant in the data, creating time varying coefficients.

variance of sales is defined for all values of $\rho \in [0,1)$, and the variance of the sales shocks, $\sigma^2_{\varepsilon_S}$, is fixed to match the variance of car sales prior to 1984 in table 1.

The relationship between production volatility and ρ in the general model, as seen in the left-most graph, is qualitatively unchanged from the first special case. Production is less volatile than sales for low values of ρ, and the ratio of production variance to sales variance increases in ρ until ρ becomes too large. Similarly, the covariance between sales and inventory investment increases in ρ, as shown in the middle graph. In the right-most graph, the variance of workforce relative to output actually declines slightly as ρ increases for values of ρ less than approximately 0.4, but the profile is relatively flat. Once ρ exceeds 0.4, the variance contribution of workforce increases in ρ just as in the second special case.

The impact of ρ on the variances is the right magnitude for explaining the changes observed in the data. Recall that the sales persistence parameter ρ decreased from 0.85 to 0.56 in the case of cars. According to the baseline parameterization of the model, this decline in ρ leads the ratio of the output variance to the sales variance to decline from well over unity to around 0.7, even greater than the change shown in Table 1. The persistence of sales for trucks declined from 0.93 to 0.69, and the model predicts a smaller decline in the variance ratio, again matching the data.

The covariance of sales with inventory investment in the model declines from 0.28 to -0.20 when ρ is reduced. In industry data this covariance declined from -0.03 in the early period to -0.20 in the late period. This simple convex approximation to the industry cost function does not replicate the exact magnitudes of all of the changes in the data, but the decline in the covariance between inventory investment and sales from the model is qualitatively consistent with industry data when the persistence of sales is reduced.

Finally, the ratio of workforce-to-output variance declines from about 0.65 to 0.50 when ρ falls. The decline predicted by the model is therefore not as dramatic as the one we observe at the plant level, where the ratio fell to 0.4. Overall, though, the predictions of the model match the data quite well.

4.4 Variance Measures, the Persistence of Sales and the Model Parameters

Could changes to parameters other than ρ have resulted in the changes measured in the auto industry data around 1984? To investigate this, the effects of various parameter values for α_1/γ_2, α_2, γ_4/γ_2 and θ on the key volatility measures (as a function of ρ) are shown in the panels of Figure 4, where alternative parameterizations appear as the dashed and dotted lines in each graph.

The effect of raising α_1/γ_2 on the volatility measures is shown in Panel A. When the penalty for deviating from the desired inventory-sales ratio is higher, the contribution of workforce adjustments to output variance moves down. The reason is that average hours per worker can adjust quickly following a sales shock, so this margin is used more intensively when maintaining inventories becomes relatively more important. The variance of output relative to the variance of sales, however, moves up when α_1/γ_2 is higher, and the correspondence between the covariance of sales with inventory investment and ρ also shifts up.

Panel B describes the effect of lowering the target inventory-sales ratio α_2, a change that has been observed in some industries (though not the automobile industry). The variance of output declines relative to the variance of sales, as expected, and inventory investment becomes less pro-cyclical. The contribution of workforce to output variance, however, increases for all values of ρ between zero and one.

Increases in the cost of adjusting workforce, γ_4/γ_2 , does move the key moments in the desired directions; the volatility of output and the covariance of sales with inventory investment move down as does the contribution of workforce to the volatility of output. The effect of an increase in γ_4/γ_2 , however, as shown in Panel C, is very small on the variance of output relative to sales and on the covariance of inventory investment with sales. Even if γ_4/γ_2 is raised to 3,000, the effects on the variance of output relative to sales and on the covariance are remarkably small.

Lastly, higher values for θ boost the variance of output relative to sales and the covariance of inventory investment and sales. This change also raises the variance contribution of workforce substantially for all values of ρ less than one because every dollar paid in workforce adjustment costs moves the minimum efficient scale of production by a larger magnitude. This, in turn, makes workforce a more cost-effective margin of adjustment.

To summarize, only two parameters, γ_4/γ_2 and ρ, are capable of moving all three key variance measures in the right direction. The effects of changes in γ_4/γ_2 on two of the key variances, however, are much smaller than the effects of changes in ρ on the variances. Thus, only changes in ρ can reproduce the patterns in the data that were documented earlier.

5 Conclusion

This paper has documented a significant change in the behavior of production, sales, and inventories in the automobile industry. The variance of production has fallen more than 70 percent since 1984, whereas the variance of sales has fallen less. The covariance of inventory investment and sales has also become more negative. Moreover, plants now rely more heavily on varying hours per worker rather than varying the number of workers.

At the same time that these changes in production scheduling occurred, the persistence of sales shocks fell significantly. Our theoretical analysis of the original Holt, Modigliani, Muth and Simon (1960) production smoothing model has shown that a reduction in sales persistence, all else equal, lowers the volatility of output relative to sales, lowers the covariance of inventory investment and sales, and reduces the portion of output volatility that stems from employment changes. Changes in the other model parameters cannot replicate all three of these changes.

Many of these changes in production volatility have been documented in other industries as well. Our analysis suggests that it would be fruitful to determine whether the behavior of sales changed in those industries as well. Determining the source of the decline in persistence is also an important area for future research.

Data Appendix

Introduction

The contribution of the variance of motor vehicle production to GDP is calculated by comparing the variances of the growth rates of the following two variables: chain-weighted total GDP and chain-weighted GDP less motor vehicles. Both of these variables are available from Table 1.2.3 from the NIPA. From 1967-2004, the variance of GDP growth is 11.30 and the variance of the growth of GDP less motor vehicles 8.65. We used nominal GDP figures to compare levels.

Figure 1 and Table 1: Car and truck sales are seasonally adjusted by the BEA and are in millions of units at an annual rate. We would have preferred to limit our analysis to light vehicles, but light-truck production is not distinguishable from total truck production prior to 1977, and data for light-truck inventories are not available before 1972. Thus, Figure 1 and Table 1 include all trucks. All production data are seasonally adjusted by the Federal Reserve Board. Car inventories are seasonally adjusted by the BEA, though we had to use our own seasonal adjustment method for trucks. For truck inventory investment, we regressed the unadjusted inventory investment on the difference between seasonally adjusted and unadjusted production and sales of trucks.

Figure 2: Car sales, car inventories, and light-truck sales are available on a seasonally-adjusted basis from the BEA. To seasonally adjust light truck inventories, we regressed the log of light truck sales on a constant and trend (allowing for breaks at 1984), as well as monthly dummy variables. We seasonally adjusted light-truck inventories by subtracting out the exponential of the fitted values of the monthly dummy variables.

Table 2: The dataset was constructed from industry trade publications in part by Bresnahan and Ramey (1994), who collected the data covering the 50 domestic car assembly plants operating in the period 1972 – 1983, and by Ramey and Vine (2004), who extended it to include all 103 car and light truck assembly plants operating in the periods 1972 – 1983 and 1990 – 2001. The data ware collected by reading the weekly production articles in *Automotive News*, which report the following variables for all North American assembly plants: (1) the number of regular hours the plant works; (2) the number of scheduled overtime hours; (3) the number of shifts operating; and (4) the number of days per week the plant is closed for (a) union holidays, (b) inventory adjustments, (c) supply disruptions, and (d) model changeovers. Observations on the line speed posted on each assembly line were collected from the *Wards Automotive Yearbook*.

Mathematical Appendix

This appendix presents the solution to our version of the Holt, Modigliani, Muth and Simon (1960) model. The derivation is based on Chapter 4 of their book, modified to allow for a discount factor and for an AR(1) sales process.

As stated in the text, the decision rules for the cost-minimization problem when sales are AR(1) take the form:

$$\begin{pmatrix} N_t \\ I_t \end{pmatrix} = C \begin{pmatrix} N_{t-1} \\ I_{t-1} \end{pmatrix} + \vec{d} \cdot S_t$$

The actual decision rules are

$$\begin{aligned} N_t \quad = \quad & \left[\frac{b_3(\lambda_1 - \lambda_2)}{\det(A)} \left(b_6 - \beta b_4 - \beta^2 \frac{b_7 - b_3(\lambda_1 + \lambda_2)}{\lambda_1 \lambda_2} \right) \right] N_{t-1} \\ & + \left[\frac{\beta^2 b_3}{\det(A)} \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 \lambda_2} \right) \right] I_{t-1} \\ & + \left[\frac{b_6 - \beta b_4 + \beta^2 b_3 \lambda_1^{-1}}{\det(A)} \left(\alpha_2 + \frac{1 - \alpha_2(1-\rho)}{1 - \lambda_2 \rho} \right) \right. \\ & \left. \quad - \frac{b_6 - \beta b_4 + \beta^2 b_3 \lambda_2^{-1}}{\det(A)} \left(\alpha_2 + \frac{1 - \alpha_2(1-\rho)}{1 - \lambda_1 \rho} \right) \right] S_t \end{aligned}$$

and

$$
\begin{aligned}
I_t \quad = \quad & \left[\frac{b_2 b_3 \left(\lambda_1 - \lambda_2\right)}{\det\left(A\right)} \left(b_6 - \beta b_4 - \beta^2 \frac{b_7 - b_3\left(\lambda_1 + \lambda_2\right)}{\lambda_1 \lambda_2} - \frac{b_1}{\beta} \right) \right. \\
& \left. - \frac{b_1\left(\lambda_1 - \lambda_2\right)}{\det\left(A\right)} \left(b_3 b_5 - b_4 b_7 + b_3 b_7 \left(\lambda_1 + \lambda_2\right) - b_3^2 \lambda_1 \lambda_2 \right) \right] N_{t-1} \\
& + \left[\frac{\beta^2 b_2 b_3}{\det\left(A\right)} \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 \lambda_2} \right) - \frac{b_1\left(\lambda_1 - \lambda_2\right)}{\det\left(A\right)} \left(b_4 - b_3\left(\lambda_1 + \lambda_2\right)\right) + 1 \right] I_{t-1} \\
& + \left[\frac{\left(b_1\left(b_5 - b_4 \lambda_2 + b_3 \lambda_2^2\right) - b_2\left(b_6 - \beta b_4 + \beta^2 b_3 \lambda_2^{-1}\right)\right)}{\det\left(A\right)} \left(\alpha_2 + \frac{1 - \alpha_2\left(1 - \rho\right)}{1 - \lambda_1 \rho} \right) \right. \\
& \left. - \frac{\left(b_1\left(b_5 - b_4 \lambda_1 + b_3 \lambda_1^2\right) - b_2\left(b_6 - \beta b_4 + \beta^2 b_3 \lambda_1^{-1}\right)\right)}{\det\left(A\right)} \left(\alpha_2 + \frac{1 - \alpha_2\left(1 - \rho\right)}{1 - \lambda_2 \rho} \right) - 1 \right] S_t,
\end{aligned}
$$

where

$$
\left[A\right] = \begin{bmatrix} b_5 - b_4 \lambda_1 + b_3 \lambda_1^2 & -\left(b_6 - \beta b_4 + \beta^2 b_3 \lambda_1^{-1}\right) \\ b_5 - b_4 \lambda_2 + b_3 \lambda_2^2 & -\left(b_6 - \beta b_4 + \beta^2 b_3 \lambda_2^{-1}\right) \end{bmatrix}
$$

and

$$
b_1 = \left(\frac{\beta \gamma_4}{\theta \gamma_2} \right)
$$

$$
b_2 = \left(\theta + b_1 \left(\frac{1 + \beta}{\beta} \right) \right)
$$

$$
b_3 = \frac{b_1 \gamma_2}{\alpha_1 \beta^2}
$$

$$b_4 = b_1\left(\frac{1}{\beta}\right) + 2b_3\left(\beta + 1\right)$$

$$b_5 = \theta + b_1\left(\frac{\beta + 1}{\beta}\right) + b_3\left(1 + 2\beta\right)$$

$$b_6 = b_1 + b_3\left(2\beta + \beta^2\right)$$

$$b_7 = \frac{b_1}{\alpha_1\beta^2}\left(\alpha_1\beta + \gamma_2\right),$$

and the relevant roots are given by[14]:

$$\lambda_i = \frac{\beta}{2}\left[\left(\phi_j + \frac{\beta+1}{\beta}\right) - \sqrt{\phi_j\left(2\frac{\beta+1}{\beta} + \phi_j\right) + \left(\frac{\beta-1}{\beta}\right)^2}\,\right] \quad \text{for } i = j = 1, 2,$$

where

$$\phi_j = \frac{b_1 \pm \sqrt{b_1^2 - 4b_3\beta^2\theta}}{2b_3\beta^2} = \frac{\alpha_1}{2\gamma_2}\left(1 \pm \sqrt{1 - \frac{4\gamma_2^2\theta^2}{\alpha_1\gamma_4\beta}}\right) \text{ for } j = 1, 2.$$

---

[14] See Holt, Modigliani, Muth and Simon, p. 100 for an explanation of why these two roots are the ones guaranteed less than one in modulus.

References

Aizcorbe, A., "Procyclical Labour Productivity, Increasing Returns to Labour and Labour Hoarding in Car Assembly Plant Employment," *Economic Journal,* 102, 1992, 860 – 873.

Automotive News, Detroit: Crain Automotive Group.

Blanchard, O. J., "The Production and Inventory Behavior of the American Automobile Industry," *The Journal of Political Economy,* 91(3), June 1983, 365 – 400.

Bresnahan, T.F. and Ramey, V.A., "Output Fluctuations at the Plant Level," *The Quarterly Journal of Economics,* August 1994, 593 – 624.

Golob, J. E., "Post-1984 Inventories Revitalize the Production-Smoothing Model," Working Paper, University of Missouri - Kansas City, 2000.

Hall, G. J., "Non-convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants," *Journal of Monetary Economics,* 45 (2000), 681 – 716.

Hetrick, R.L., “Analyzing the Recent upward Surge in Overtime Hours,” *Monthly Labor Review*, February 2000, 30 – 33.

Holt, C.C., Modigliani, F., Muth, J.F., and Simon, H.A., *Planning Production, Inventories and Workforce,* Englewood Cliffs, N.J.: Prentice-Hall, 1960.

Kahn, J. A., McConnell, M. M. and Perez-Quiros, G., "On the Causes of the Increased Stability of the U.S. Economy," *FRBNY Economic Policy Review*, May 2002, 183 – 202.

Kim, C. and Nelson, C. R., "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," The *Review of Economics and Statistics,* 81(4), November 1999, 608 – 616.

McConnell, M.M., and Perez-Quiros, G., "Output Fluctuations in the United States: What has Changed Since the Early 1980s?" *American Economic Review*, 90(5), December 2000, 1464 – 76.

Ramey, V. A. and Vine, D. J., “Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry,” NBER working paper 10384, March 2004.

Ramey, V. A. and West, K. D., "Inventories," in *Handbook of Macroeconomics,* eds. John D. Taylor and Michael Woodford, Elsevier Science Press, 1999.

Table 1: Decomposition of Motor Vehicle Output Volatility

(*Physical units, cars and trucks separately*)

	Cars		Trucks	
	1967:2-1983:12	1984:2-2004:12	1967:2-1983:12	1984:2-2004:12
$Var(Y)$	3.13	0.92	9.16	1.15
$Var(S)$	2.45	1.01	6.89	0.94
$Var(\Delta I)$	0.20	0.15	0.29	0.09
$Cov(S,\Delta I)$	-0.03	-0.20	0.15	-0.08
$\frac{Var(Y)}{Var(S)}$	1.28	0.91	1.33	1.22

Y = production, S = sales, and ΔI = change in dealer inventories. Production and sales were normalized by the exponential of a fitted linear trend to the log of the variable, estimated separately over each period. Inventory investment was normalized by the fitted trend in the log level of inventories. The variances and covariances in the table are 100 times the actual ones. (See data appendix for data sources and details.) Data were seasonally adjusted.

Addendum: Comparison of Production and Sales Including Imports and Exports

	Cars	
	1967:2-1983:12	1984:2-2004:12
$Var(\text{U.S. Production + N. American Imports})$	2.94	0.94
$Var(\text{U.S. Sales + Exports})$	2.40	0.96
$Var(\Delta I)$	0.20	0.15
$Cov(S,\Delta I)$	-0.01	-0.17
$\frac{Var(Y)}{Var(S)}$	1.23	0.98

Table 2: Importance of Intensive and Extensive Margins of Adjustment for the Variance of Monthly Motor Vehicle Output *

(Percent of average plant-level variance attributed to use of each margin)

Margin of Adjustment	1972 – 1983	1990 – 2001
Changes in Extensive Margins	**64.7 %**	**37.3 %**
Shifts	*33.0*	*5.4*
Line Speeds	*21.1*	*17.8*
Changes in Intensive Margins	**39.6 %**	**51.0 %**
Temporary Closures (Inventory Adjustments)	*32.5*	*33.1*
Overtime Hours	*9.6*	*24.3*

* Plant-level variance is calculated after holidays, supply disruptions, model changeovers and extended closures are removed. Percent impact of each margin on output variance is calculated by comparing variance of actual production with the variance of hypothetical production if each margin (in turn) were held fixed. Contributions to variance are the weighted average among all plants operating in each period. Contributions of extensive and intensive margins do not sum to 100 because of covariance terms. The same is true for individual margins within each category. See data appendix for data sources.

Table 3: Estimates of Aggregate Automobile Sales Process

Coefficients (and standard errors) from the regression:

$$\tilde{S}_t = \alpha_0 + \alpha_1 \cdot t + \alpha_2 \cdot \tilde{S}_{t-1} + D_t \cdot \left[\beta_0 + \beta_1 \cdot t + \beta_2 \cdot \tilde{S}_{t-1} \right] + \varepsilon_t$$

$$\text{where } \varepsilon_t \sim N\left(0, \sigma^2 + \beta_3 D_t\right)$$

$$\text{and } \begin{array}{l} D_t = 0 \text{ for } t < 1984:1 \\ D_t = 1 \text{ for } t \geq 1984:1 \end{array}$$

Coefficient		Cars	Light Trucks	Light Vehicles
α_0	*(constant)*	0.367** (.136)	-00045 (.042)	0.279** (.117)
β_0	*(Δ constant)*	0.628** (.215)	0.287** (.086)	0.654** (.234)
α_1	*(trend)*	-0.0002 (.00013)	0.00017 (.00016)	-.00005 (.00011)
β_1	*(Δ trend)*	-0.0003** (.00017)	0.0011** (.00035)	0.00050** (.00020)
α_2	*(AR(1))*	0.851** (.049)	0.934** (.030)	0.886** (.043)
β_2	*(Δ AR(1))*	-.293** (.108)	-.243** (.076)	-.271** (.102)
σ^2	*(innov. Variance)*	0.0070** (.00116)	0.0089** (.0012)	0.0066** (.0011)
β_3	*(Δ innov. Variance)*	-0.00046 (.0016)	-.0040** (.0014)	-.0013 (.0014)
Log likelihood		490.3	502.3	523.2

Standard errors were computed using Eicker-White methods.
** denotes significant at the 5 % level.
Sample is 1967:2 – 2004:12, N = 455

Panel A: Cars

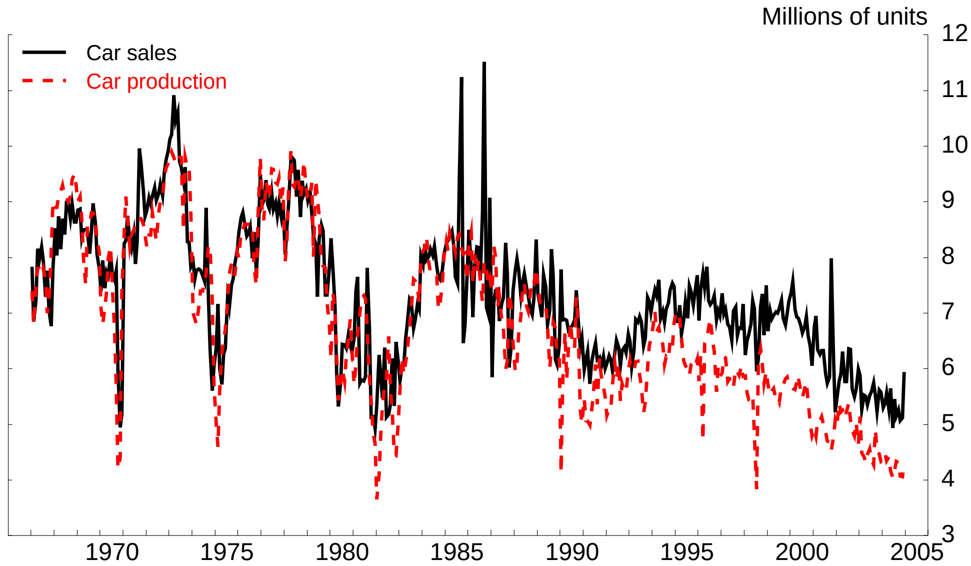

Panel B: Trucks

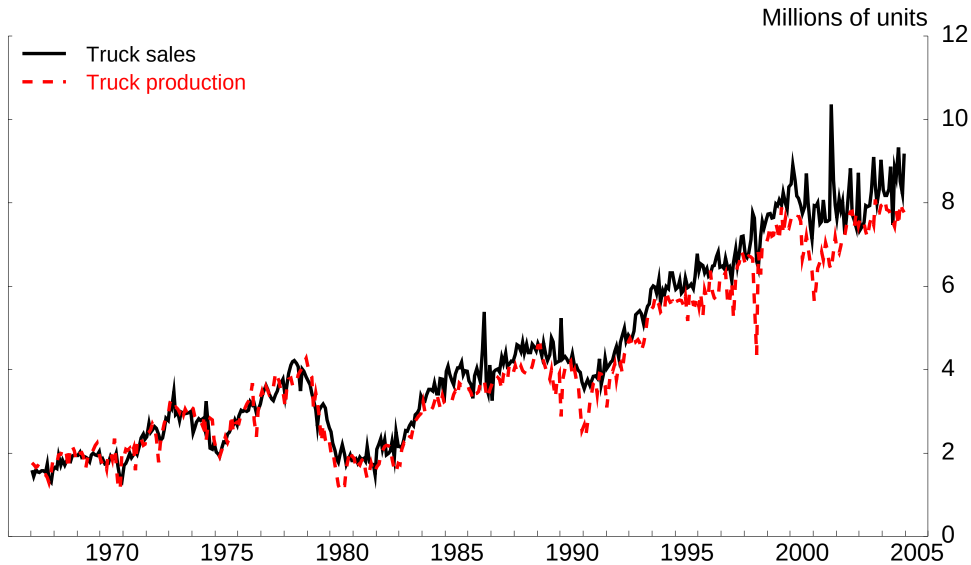

Figure 1: U.S. Automobile Production and Domestic Sales*

January 1967 to December 2004

(Physical Units, Annualized Basis)

* Domestic Sales include vehicles built in the U.S., Canada and Mexico

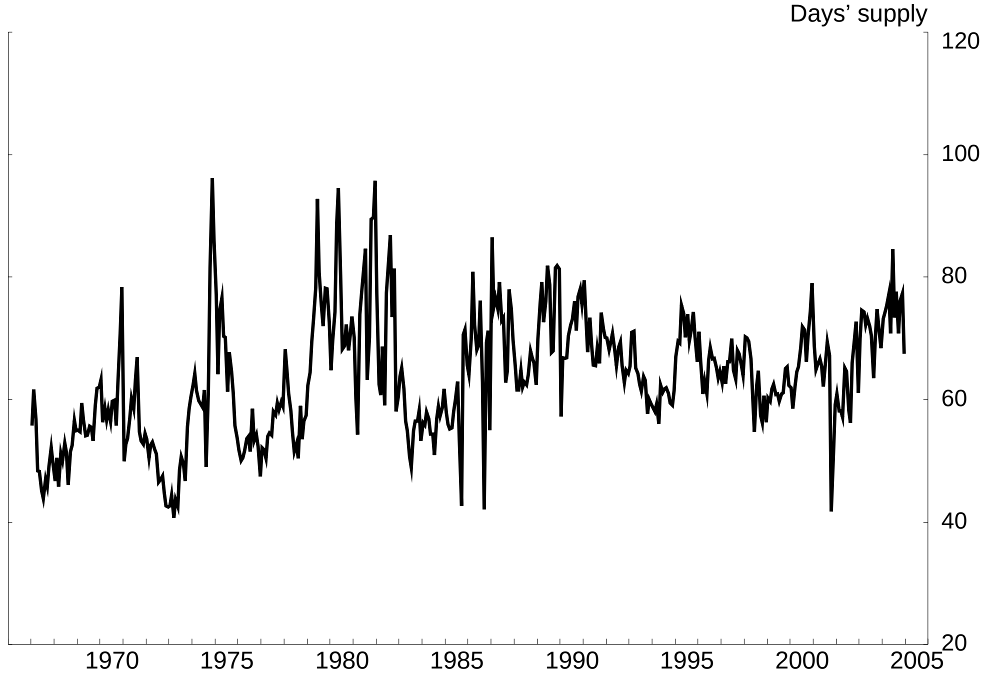

Figure 2: Inventory to Sales Ratio for Domestic Cars and Light Trucks

January 1967 to December 2004

(Days of physical units**)**

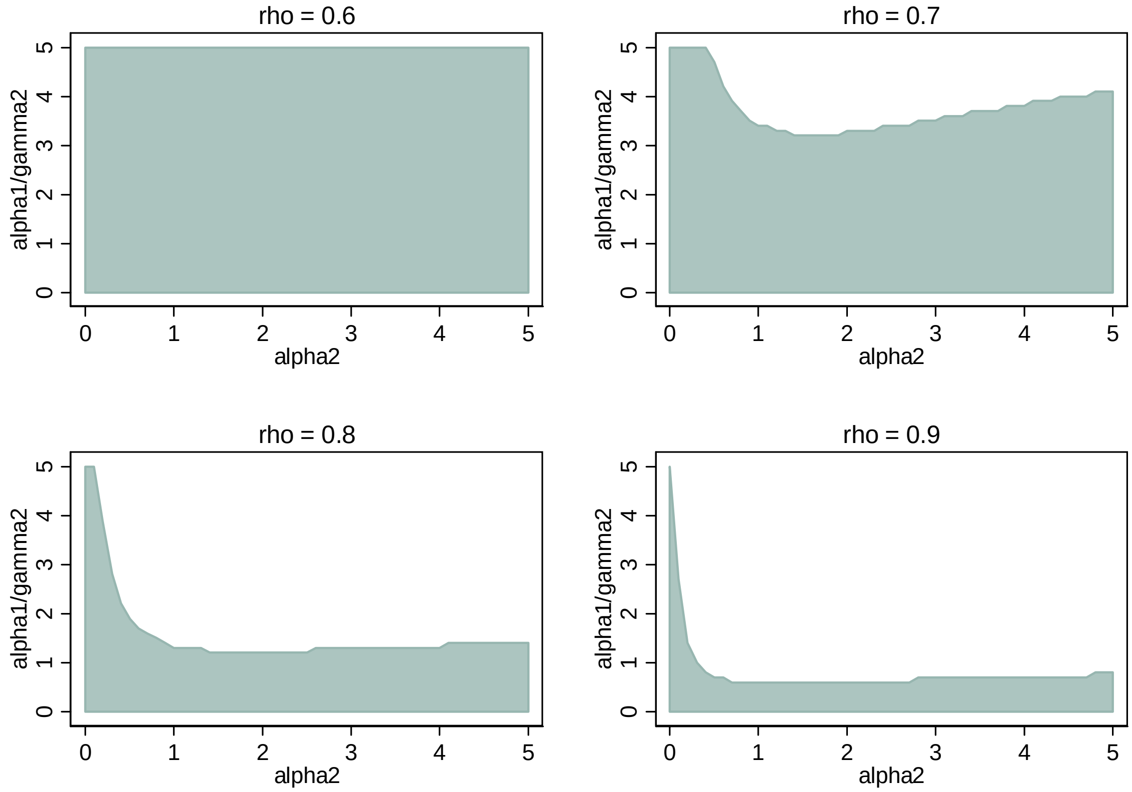

Figure 3: Parameter Regions for which $\frac{\partial\left(Var(Y)_t / Var(S_t)\right)}{\partial\rho} > 0$

Figure 4: Variance of Sales, Output and Employment as a Function of ρ *

Panel A: Effect of (α_1/γ_2)

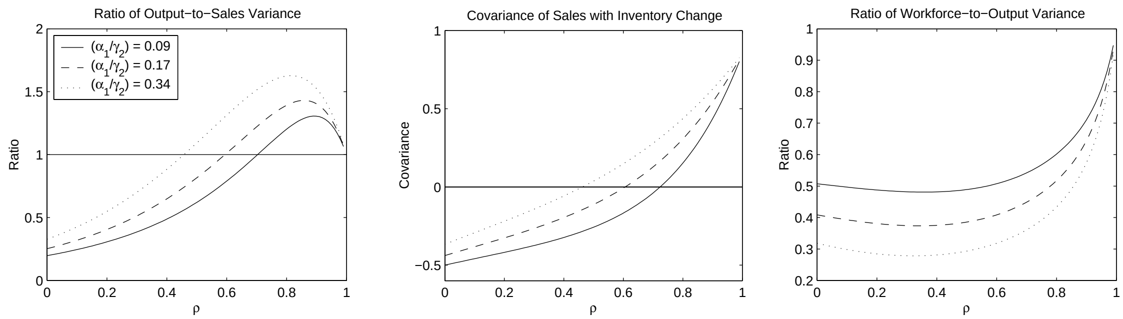

Panel B: Effect of α_2

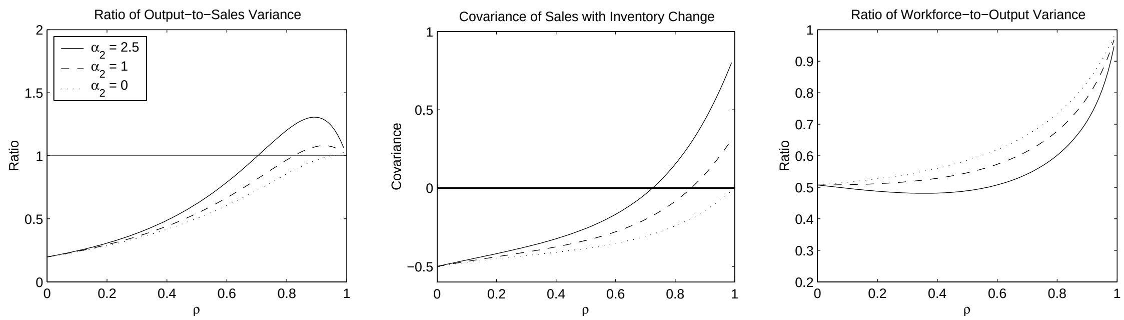

Panel C: Effect of (γ_4/γ_2)

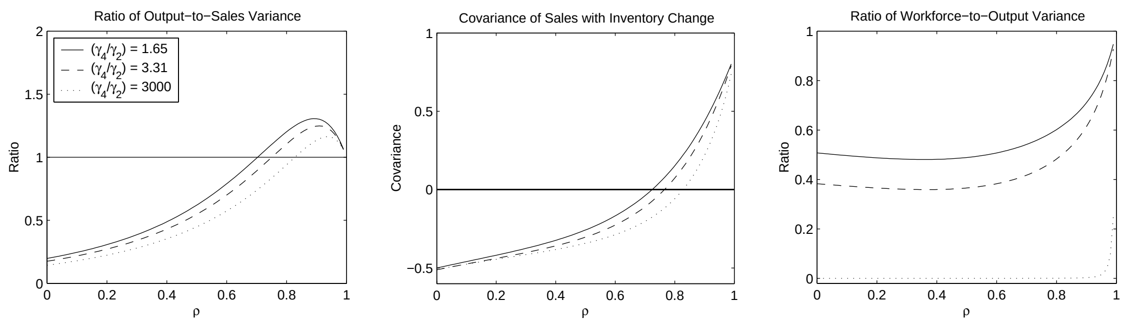

Panel D: Effect of θ

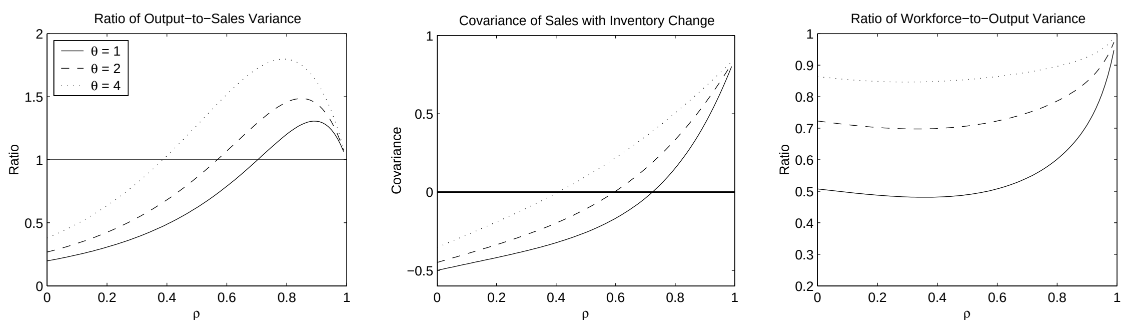

* Parameters not subject to variation in each graph are fixed at their benchmark level: $\alpha_1/\gamma_2 = .07$, $\gamma_4/\gamma_2 = 1.8$, $\theta = 1$, $\alpha_2 = 2.5$, $\beta = 0.997$. Variance of sales innovations is fixed so that the variance of sales equals 2.45 when $\rho = .85$. This implies that covariance is measured on the same scale as cars in Table 1.